



Council of European National
Top-Level Domain Registries



IETF 124

The **124rd IETF meeting** took place in Montreal, Canada between 1 and 7 November 2025 with over 150 sessions, a 2-day hackathon, and a wide range of side events. Pawel Kowalik and Christian Simmen from DENIC attended the meeting and have written a summary of the main points of relevance for the CENTR members.

PAWEL KOWALIK & CHRISTIAN SIMMEN, DENIC EG





Table of contents

Introduction **3**

Informal activities **3**

Hackathon 3

IEPG 4

Formal proceedings **5**

Authenticated Transfer Protocol (ATP) BoF 5

DNSOP WG 6

DELEG WG 8

REGEXT WG 9

RPP WG 10

DCONN WG 11

Epilogue **13**

Introduction

The mission of the Internet Engineering Task Force (IETF) is to make the internet better. Most of the IETF's work is done online, but the organisation also holds 3 meetings a year. The 124th IETF meeting was held in Montreal, Canada, from 1 to 7 November 2025.

With **1,625 registered participants**, the meeting saw strong attendance, with 942 participants (58%) present on site and 683 participating remotely via the Meetecho platform. The meeting was hosted by Comcast and NBCUniversal at the Fairmont The Queen Elizabeth hotel.

The **IETF Hackathon** in the weekend prior to the meeting attracted significant participation, with work continuing on numerous projects spanning DNS, post-quantum cryptography, and domain registration protocols. There was also a '**code sprint**', at which a small group of volunteers worked at improving the tools made available by the IETF, such as the well-known **Datatracker**.

Every IETF meeting has a **packed programme**. The latest week-long gathering featured over 150 working group sessions, the **IEPG meeting**, **HotRFC Lightning Talks**, a **plenary session**, and a wide variety of side meetings.

Informal activities

Hackathon

The now traditional Hackathon got under way on the Saturday prior to the main proceedings. At the Hackathon, the applicability and interoperability of new concepts were tested on a collaborative, non-competitive basis. Groups were formed spontaneously to work on a **range of experiments** from a list of over 60 projects. A total of 682 participants interacted and programmed together intensively throughout the weekend. The Hackathon concluded with result **presentations**.

RESTFUL PROVISIONING PROTOCOL (RPP)

The **RPP hackathon** continued its goal of prototyping an alternative to EPP for domain name provisioning. The team worked on experimental implementations of the RPP core protocol and explored alternative representations for DNS data.

DNS PROJECTS

Traditionally the "DNS table" was **full of projects** focused on DNS operations and performance. The development focused on implementations of updates coming from -05 version of DELEG specification, DNS transport signalling as well as the concept of sharing DNS resolver cache for better performance and scalability.

POST-QUANTUM CRYPTOGRAPHY ACTIVITIES

Continuing the trend from IETF 122 and 123, there was significant activity around post-quantum cryptography. The [PQC DNSSEC MTL Mode Update](#) project, following the discussion and the momentum after [PQC Side Meeting at IETF 123](#), continued various experiments and interoperability tests around PQC implementations.

Reference Open-Source	Link	Algorithms
MTL reference library	https://github.com/verisign/MTL	MTL mode with SLH-DSA and ML-DSA
MTL LDNS library	https://github.com/verisign/mtl-mode-ldns	RSA, ECDSA, ML-DSA ⁽¹⁾ , FL-DSA ⁽¹⁾ , SLH-DSA ⁽¹⁾ , Mayo I/II ⁽¹⁾ , SQI Sign ⁽¹⁾ , Hawk ⁽¹⁾ , SNOVA ⁽¹⁾ , SLH-DSA w/MTL mode ⁽¹⁾ , ML-DSA w/MTL mode ⁽¹⁾
NSD [authoritative name server]	https://github.com/NLnetLabs/nsd/pull/397	RSA ⁽²⁾ , ECDSA ⁽²⁾ , ML-DSA ⁽²⁾ , FL-DSA ⁽²⁾ , SLH-DSA ⁽²⁾ , Mayo I/II ⁽²⁾ , SQI Sign ⁽²⁾ , Hawk ⁽²⁾ , SNOVA ⁽²⁾ , SLH-DSA w/MTL mode ⁽²⁾⁽³⁾ , ML-DSA w/MTL mode ⁽²⁾⁽³⁾
Unbound [recursive resolver]	https://github.com/verisign/mtl-mode-unbound	RSA ⁽²⁾ , ECDSA ⁽²⁾ , ML-DSA ⁽²⁾ , FL-DSA ⁽²⁾ , SLH-DSA ⁽²⁾ , Mayo I/II ⁽²⁾ , SQI Sign ⁽²⁾ , Hawk ⁽²⁾ , SNOVA ⁽²⁾ , SLH-DSA w/MTL mode ⁽²⁾⁽³⁾ , ML-DSA w/MTL mode ⁽²⁾⁽³⁾
BIND [authoritative and recursive resolver]	TBD	RSA, ECDSA, FL-DSA ⁽¹⁾ , Mayo ⁽¹⁾ , SQI Sign ⁽¹⁾ , Hawk ⁽¹⁾ , ANTRAG-512, SLH-DSA w/MTL mode ⁽¹⁾
Core DNS [authoritative]	https://github.com/fjblanco/mtl_coredns_plugin	RSA, ECDSA, ML-DSA ⁽¹⁾ , FL-DSA ⁽¹⁾ , SLH-DSA ⁽¹⁾ , Mayo I/II ⁽¹⁾ , SNOVA ⁽¹⁾ , SLH-DSA w/MTL mode ⁽¹⁾

Reference Source	Link	Algorithms
Core DNS [authoritative]	https://github.com/fjblanco/mtl_coredns_plugin	RSA, ECDSA, ML-DSA ⁽¹⁾ , FL-DSA ⁽¹⁾ , SLH-DSA ⁽¹⁾ , Mayo I/II ⁽¹⁾ , SNOVA ⁽¹⁾ , SLH-DSA w/MTL mode ⁽¹⁾

1 - Enabled at compile time, depends on additional cryptographic libraries
 2 - When signed with LDNS.
 3 - Includes POC for MTL mode EDNS option.

4

Figure 1 PQC Name server implementations

IEPG

The Internet Engineering and Planning Group (IEPG) held its traditional Sunday pre-IETF meeting featuring operational topics including DNS, routing security, and IPv6 deployment measurements.

- [Geoff Huston](#) presented on making local root zones the default for DNS resolvers. With root server query load growing 40% in two years to 130 billion queries daily, the presentation advocated for recursive resolvers to maintain a complete local copy of the root zone file (2.2 MB), validated using the Zone MD record. This approach improves query performance, enhances privacy by eliminating unnecessary root queries, and reduces dependency on root server infrastructure. Three major recursive resolver implementations (BIND, Knot, Unbound) already support local root zone operation, positioning it as a practical deployment option.
- [Shane Kerr](#) examined IPv4 versus IPv6 performance for authoritative DNS servers using RIPE Atlas measurements. Results showed IPv6 query performance matching or slightly exceeding IPv4 in well-connected networks.
- [Tobias Fiebig](#) explored DNS resolution through NAT64 environments, benchmarking open-source NAT64 implementations and measuring their impact on DNS query timeouts.

- **Job Snijders** proposed ERIC synchronization, a next-generation RPKI data distribution protocol addressing scalability limitations in existing RSYNC and RRDP approaches.

Formal proceedings

A few of the many working group sessions are outlined below. Because many of the sessions take place in parallel, it can be difficult to choose which one(s) to attend - therefore the list might not cover all the topics from all the working groups, it is the author's choice of relevant ones to the CENTR community.

Authenticated Transfer Protocol (ATP) BoF

The Authenticated Transfer Protocol Birds of a Feather session explored standardisation opportunities for a new social networking protocol framework developed by Bluesky. ATP addresses fundamental challenges in federated social media: enabling global aggregation and redistribution of content across organisational boundaries, supporting individual account portability between service providers, and separating identity control from network location. The framework employs self-certifying public data repositories, content-addressed data structures, and an abstracted identity layer that allows users to move between services while maintaining their social graph and content history.

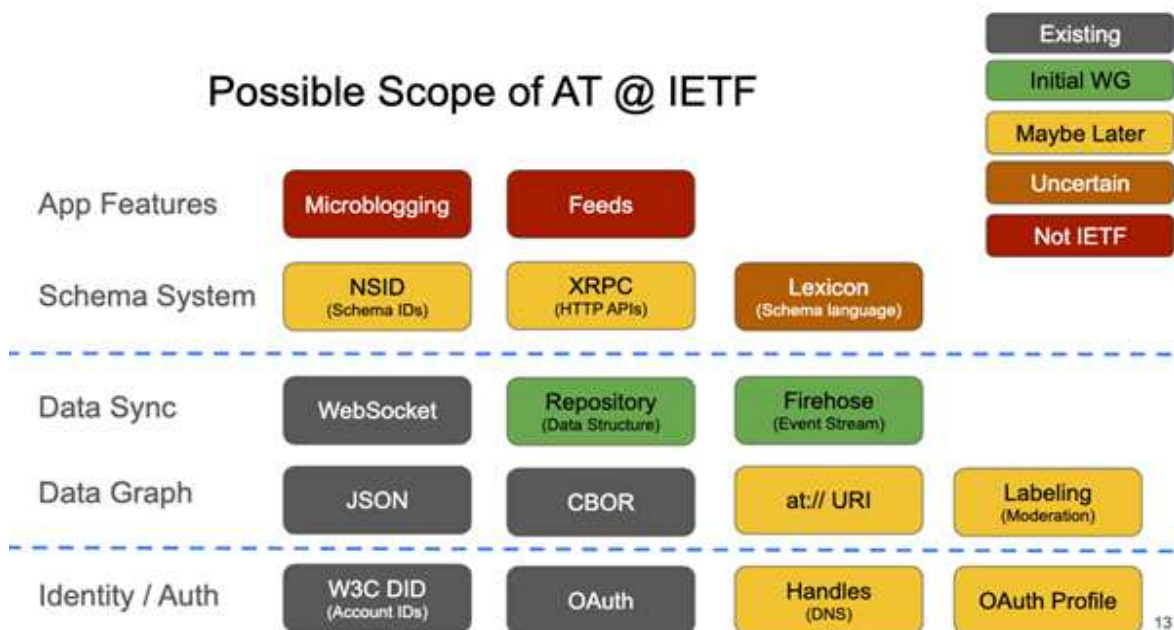


Figure 2 Building blocks of AT protocol

The session attracted significant participation from the AT ecosystem including developers from Bluesky (40 million users), Flipboard, Tangled (social coding platform), and Blacksky (self-governing communities), alongside strong support from IETF community members.

Current deployments demonstrate diverse applications beyond microblogging, including custom algorithmic feeds (over 50,000 third-party feeds on Bluesky), community-driven content moderation services, social coding platforms, and self-governing online communities.

DOMAIN NAMES AS IDENTITY HANDLES

A key architectural element of ATP relevant to domain name operators is its use of domain names as user identity handles. The protocol employs an abstracted identity layer that separates account identity from network location, enabling account portability between social media providers while maintaining a persistent, verifiable identity.

Users can claim domain names as their handles (e.g., @example.com), providing human-readable identifiers that are cryptographically verifiable and portable across ATP-based services. Domain ownership proves identity authenticity, like how certificate authorities verify domain control for TLS certificates. This approach enables users to maintain their identity when switching between ATP service providers, with the domain name serving as the stable identifier. Organisations can use their domain names for authenticated presence across federated social networks, maintaining brand consistency while retaining decentralised control over their identity rather than depending on centralised social media platforms.

The protocol uses DIDs (Decentralised Identifiers) published over DNS for the underlying identity layer, with domain names providing a user-facing handle that maps to cryptographic keys. This architecture creates opportunities for registries and registrars to offer identity services complementing traditional domain registration, potentially including verification services, key management, and identity portability features.

The BoF, even though officially not yet working group forming, concluded with strong community support for future forming a working group. The proposed scope includes data structures, synchronisation protocols, and public data handling, while deliberately excluding application-level features and leaving the identity layer (including domain-based handles) for separate consideration. The IETF community emphasised that ATP provides useful building blocks for internet-scale authenticated content networks.

DNSOP WG

The DNSOP (DNS Operations) Working Group addresses operational aspects of the DNS protocol. The working group held two sessions at IETF 124 covering IPv6 transport, domain verification, cryptographic algorithm management, and delegation enhancements. A few highlights are presented below. For a list of all discussed drafts can be found on the [meeting agenda](#).

PUBLISHED WORK

Since the previous IETF meeting the following RFC have been published:

- [RFC9824 - Compact Denial of Existence in DNSSEC](#) - Enables online signing servers to

generate compact NXDOMAIN responses using a single NSEC/NSEC3 record, reducing signing operations and response sizes while preventing zone enumeration attacks. This optimisation benefits operators using dynamic DNSSEC signing. This document updates RFCs [4034](#) and [4035](#).

- [RFC9859 - Generalised DNS Notifications](#) - This document extends DNS NOTIFY ([RFC 1996](#)) beyond zone transfers to trigger automated actions such as parental DS record updates and DNSSEC bootstrapping. Introduces the DSYNC record type for discovering notification endpoints, enabling timely propagation of delegation changes without polling.

WORK IN PROGRESS

DNS IPv6 Transport Operational Guidelines

Momoka Yamamoto presented [draft-ietf-dnsop-3901bis](#), updating operational guidelines to make IPv6 support mandatory for DNS infrastructure. The document addresses the reality that IPv4 quality is degrading in some regions while IPv6 deployment continues to grow. By establishing IPv6 as a baseline requirement rather than optional, the guidelines will help operators ensure reliable DNS service as the internet transitions away from IPv4 dependency. The draft received strong support from implementers who noted it provides essential documentation for justifying IPv6 deployment to management and users. Discussion continues on specific recommendation levels for different resolver types.

POTENTIAL NEW WORK

Phasing In and Out of Cryptographic Algorithms

Two drafts address the critical challenge of managing DNSSEC cryptographic algorithm transitions. As algorithms age or vulnerabilities emerge, the DNS community needs systematic approaches for deprecating old algorithms and introducing new ones without breaking DNSSEC validation.

- [draft-crocker-dnsop-dnssec-algorithm-lifecycle](#) proposes a framework for documenting algorithm lifecycle stages—from introduction through deprecation to prohibited status. This would provide clear guidance for zone operators and validator implementers on when algorithms should be deployed, when they should be phased out, and how to communicate these transitions to the community. The work recognises that DNSSEC's global deployment and long-lived signed records create unique requirements distinct from other protocols like TLS.
- [draft-huque-dnsop-multi-alg-rules](#) addresses the practical mechanics of running multiple algorithms simultaneously, particularly for multi-signer deployments where different providers may use different algorithms, and during algorithm rollovers. The draft proposes relaxed validation rules that prevent zones from becoming insecure when validators cannot process all advertised algorithms, enabling graceful degradation rather than complete

Provision for Range of RRTYPEs in Parent

Roy Arends presented [draft-arends-dnsop-delext](#) proposing to reserve a range of DNS resource record types specifically for delegation-related records in parent zones. Currently, each new delegation-related RRtype requires individual allocation from the general registry. By establishing a dedicated range, future delegation enhancements (such as DELEG and related records) can be deployed more systematically. The proposal also carves out private use space for experimentation, facilitating development and testing of new delegation mechanisms before standardisation.

For-Sale TXT Record

Marco Davids presented [draft-davids-forsalereg](#), proposing a standardised “_for-sale” DNS name for registries to signal domain name availability for purchase. This addresses a common operational need: providing a machine-readable, globally consistent way to indicate domains are for sale, augmenting inconsistent web-based parking pages and proprietary signalling methods. The standardised approach would benefit domain investors, registrars, and automated domain acquisition systems by providing a reliable discovery mechanism. Given the limited scope and primary relevance to registry operations rather than protocol work, the working group suggested pursuing publication through the Independent Submission Editor rather than as a working group item.

DELEG WG

The DELEG (DNS Delegation) Working Group is developing a modern replacement for the NS record to enable richer delegation information in the DNS. The working group is progressing toward completion of its base protocol specification with multiple implementations demonstrating deployment readiness.

IMPLEMENTATION MILESTONE

The working group reached a significant milestone with four independent implementations providing feedback: Akamai (authoritative and recursive), BIND (dig client and zone loading), CoreDNS (RR format), and dnspython (RR format). Implementation experience revealed that nested multi-level `include-delegi=` indirection proved awkward, though necessary for operational flexibility. Registry outreach produced positive results with no fundamental objections from TLD registries and RIRs.

BASE PROTOCOL AND MAJOR DECISIONS

[draft-ietf-deleg](#) progressed from version -01 to -05, maintaining core functionality while adding test vectors, expanded examples, and RESINFO integration. The working group will develop comprehensive test vectors to ensure implementation interoperability.

The working group decided to publish the requirements document alongside the base pro-

protocol for historical context, ensuring future implementers understand design rationale. The group committed to acknowledging which requirements were actually met rather than retroactively modifying requirements.

Strong support emerged for rechartering to incorporate proposed extensions (currently out of scope) sooner than originally anticipated, enabling parallel work on transport extensions and other enhancements that have been waiting on base protocol completion.

REGEXT WG

The REGEXT (Registry Extensibility) working group is focused on coordinating and developing standard extensions for key Internet registry protocols. Its primary role is to manage the standardisation process for extensions to the Extensible Provisioning Protocol (EPP), used for domain name registration, and the Registration Data Access Protocol (RDAP), used for retrieving registration data.

PUBLISHED WORK

Since IETF 123, three RFCs were published:

- [RFC 9874 - Best Practices for Deletion of Domain and Host Objects in EPP](#) - Clarifies delete operation semantics in EPP, helping registries and registrars align implementations with consistent behaviour
- [RFC 9873 - Additional Email Address Extension for EPP](#) - Enables internationalised email addresses in EPP, supporting global users with non-ASCII email addresses
- [RFC 9877 - RDAP Extension for Geofeed Data](#) - Provides geolocation data access via RDAP, useful for network operators managing IP address geolocation information

WORK IN PROGRESS

Updates were provided on several ongoing drafts.

RDAP Extension for TTL Values

Gavin Brown presented [draft-ietf-regext-rdap-ttl-extension](#), enabling registries to expose DNS Time-To-Live values through RDAP responses. This allows registrars and DNS operators to programmatically discover TTL values for delegation records (NS, DS, A, AAAA), which is particularly valuable during operational changes like DNS server migrations or DNSSEC key rollovers where knowing TTL values helps operators calculate safe timing for changes. The extension complements the EPP TTL extension (RFC 9803) by providing read access to the values that can be modified via EPP. The working group is finalising the JSON representation and will proceed toward working group last call.

RDAP Referrals

Gavin Brown presented [draft-ietf-regext-rdap-referrals](#), proposing an efficient mechanism for RDAP servers to redirect clients to the authoritative RDAP server for a domain. This solves a practical problem: when querying a registry RDAP server for a domain, clients currently receive the full domain object even when they only need to discover which registrar operates it. The referral mechanism reduces bandwidth and processing overhead by providing just the redirect information, particularly beneficial for thin registry models where the registry holds minimal data and the registrar holds complete registration details. The working group favoured the path segment redirect approach as the most client-friendly solution.

POTENTIAL NEW WORK

Domain Variant Support for EPP

Jim Galvin presented a revised version of [draft-galvin-regext-epp-variants](#), proposing a mechanism to manage groups of related domain names with coupled lifecycles. This addresses operational needs for bundling domain names—whether IDN variants, diacritic variants, or other related groupings—where creating, transferring, or deleting one domain affects the entire group. The proposal enables single EPP commands to manage the group atomically, preventing inconsistent states and reducing operational complexity for both registries and registrars. The draft shifted from “variant” terminology to “related groups” to reflect wider use-case of this extension.

Balance Mapping for EPP

James Gould presented [draft-gould-regext-balance](#), proposing a standardised extension for registrars to monitor their account balance and available credit at registries. By consolidating proprietary and fragmented approaches like [Low Balance Mapping](#) or [Balance Mapping](#) into a standard mechanism, registrars can implement consistent monitoring across multiple registries. The working group will consider adoption.

RPP WG

The RESTful Provisioning Protocol (RPP) Working Group is developing a standardised protocol for domain name management between registries and registrars, offering a modern REST- and JSON-based alternative to the existing Extensible Provisioning Protocol (EPP).

The working group is focused on achieving consensus on the [requirements document](#) before IETF 125. This milestone is critical for advancing to the core architecture and extension mechanism deliverables.

WORK IN PROGRESS

Requirements

Maarten Wullink presented updates to [draft-ietf-rpp-requirements](#). The -02 revision resolved numerous issues from IETF 123, including decisions on strict data validation, response caching, bless requests/responses, and transaction information in headers. Two interconnected issues remain open: client data omission and mandatory profile support, requiring working group input.

The authors emphasised the draft is ready for review and strongly encouraged community members to provide feedback now rather than waiting for the consensus call. The working group aims to achieve consensus before IETF 125. Once consensus is reached, the document will be placed in “Parked” status (allowing updates during RPP development with renewed consensus) and eventually published as an Informational RFC when RPP design and development work is completed.

Design Work

Progress continues on design documents that will implement the requirements. [Draft-kowalik-rpp-architecture](#) describes RPP’s layered architectural principles, mapping HTTP transport, data elements, and operations layers to specific implementation documents. [Draft-wullink-rpp-core](#) explores concrete protocol mechanisms including code headers, authorisation, process handling, and availability checks. [Draft-simmen-rpp-dns-data](#) proposes JSON representation for DNS resource records, with active discussion on delegation control and DELEG integration. These design documents will be aligned with finalised requirements and proceed toward adoption.

Pawel Kowalik presented [draft-kowalik-rpp-data-objects](#), defining data elements and operations in an abstract, representation-agnostic model. The draft covers primitive data types, associations, component objects, and four core data object types: domain, contact, host, and organisation. An IANA RPP Data Object Registry is proposed for discoverability. The draft is work-in-progress seeking early feedback.

DCONN WG

The Domain Connect Working Group held its inaugural session at IETF 124, becoming one of the most recently formed working groups. DCONN focuses on standardising an existing, widely deployed protocol that automates DNS configuration for third-party services.

THE PROBLEM AND SOLUTION

Setting up third-party services like Microsoft 365 or Google Workspace requires configuring multiple DNS records—a complex task that often fails even with extensive documentation. Microsoft 365, for example, requires 7-15 DNS records in a six-step process, achieving only 50% success rates despite providing 40 minutes of training materials and 16 help sites.

Domain Connect solves this through automated DNS provisioning using templates. Service providers define required DNS records in templates that DNS operators pre-deploy and vouch for. When users activate a service, they authenticate with their DNS provider, review proposed changes, and authorise configuration with a single click—reducing hours of manual work to seconds.

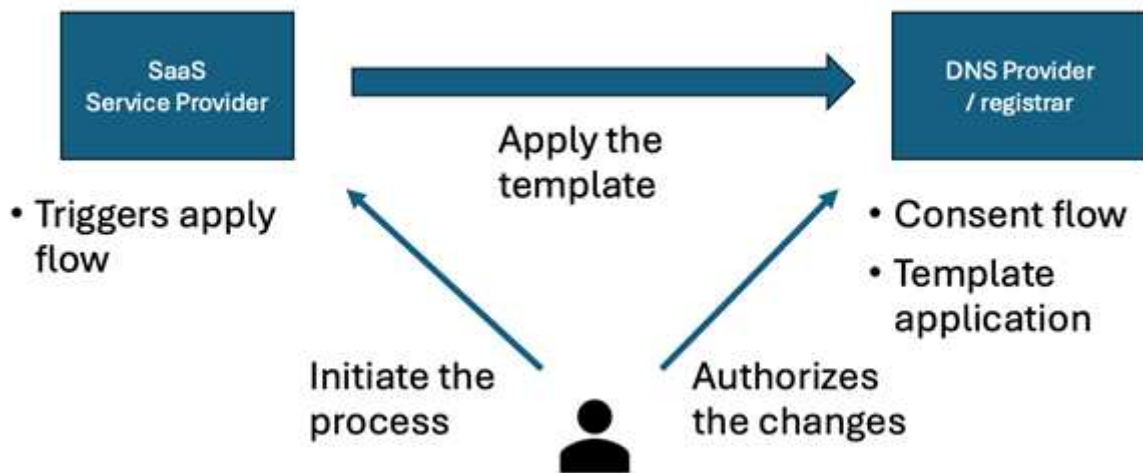


Figure 3 Domain Connect protocol flow

The protocol is already deployed at approximately 20 DNS providers and hundreds of service providers, demonstrating proven operational value. Users discover DNS provider support via `_domainconnect` TXT records, service providers redirect users to DNS provider authorisation endpoints, and configurations are applied automatically after user consent.

WORKING GROUP SCOPE

The working group maintains a deliberately narrow scope: standardising the existing protocol to ensure specification precision, address ambiguities, and resolve security considerations - not developing new features. This focused approach reflects the protocol's journey from initial IETF presentation in 2016 through independent development and deployment, returning to IETF in 2024 for formal standardisation.

For domain registries and registrars, Domain Connect reduces support burden for DNS configuration issues, improves customer experience through seamless service integration, and provides competitive differentiation through one-click setup for major service providers. IETF standardisation ensures interoperability and provides security review as deployment expands.

The working group showed strong support to adopt [draft-kowalik-domainconnect](#) as a starting point of its effort.

Epilogue

IETF 124 was another successful gathering, for which the hosts Comcast and NBCUniversal deserve congratulations. There was once again a full and stimulating program, and ample opportunity for participants to network. The focus is firmly on the future, within both the IETF working groups and the IRTF working groups. And, because the future is exciting, IETF meetings are consistently stimulating and relevant.

The **next IETF meeting** is scheduled for 14 to 20 March 2026 in Shenzhen, China.



**Council of European National
Top-Level Domain Registries**



About CENTR

CENTR is the association of European country code top-level domain (ccTLD) registries, such as .de for Germany or .si for Slovenia. CENTR currently counts 51 full and 10 associate members – together, they are responsible for over 80% of all registered domain names worldwide.

The objectives of CENTR are to promote and participate in the development of high standards and best practices among ccTLD registries.

Full membership is open to organisations, corporate bodies or individuals that operate a country code top level domain registry.

CONTACT

 **CENTR VZW/ASBL**
Belliardstraat 20
1040 Brussels, Belgium
0885.419.166 | RPR Brussels

 +32 2 627 5550

 secretariat@centr.org

 www.centr.org

FOLLOW US

To keep up-to-date with CENTR activities and reports, follow us on LinkedIn



© This publication has been authored by CENTR. Reproduction of the texts of this publication is authorised, provided the source is acknowledged.